ATLAS HL-LHC Demonstrators with Data Carousel: Data-on-Demand and Tape Smart Writing

M. Borodin, A. Klimentov, T. Korchuganova, M. Lassnig, J. Leduc, T. Maeno,
H. Musheghyan, D. Ressmann and **X. Zhao**

On behalf of the ATLAS Computing Activity

CHEP, 2024
Krakow, Poland

Outline

- ATLAS HL-LHC demonstrators with Data Carousel
- Data on demand demonstrator
- Tape smart writing demonstrator
- Conclusion and next steps

Team Effort ---

- *workflow management system team(WFMS)*
- *distributed data management team (DDM/Rucio)*
- *distributed production and analysis team (DPAs)*
- *operations team(Ops)*
- *monitoring team*
- *ATLAS distributed computing(ADC) coordinators and experts*
- *CERN T0 and all T1s storage and tape experts.*

ATLAS HL-LHC demonstrators with Data Carousel

- ATLAS HL-LHC demonstrators
  - ATLAS HL-LHC demonstrator projects are initiatives designed to test and validate new technologies and solutions across various areas, including the ATLAS software and computing, to meet the challenges of the HL-LHC upgrade. Results from these projects will be incorporated into the ATLAS HL-LHC Technical Design Report (TDR).
  - In this talk we will report on progress and results from two HL-LHC demonstrators that are from the ATLAS Data Carousel activity : Data-on-demand and Tape smart writing.

- ATLAS Data Carousel
  - A project started in June 2018, to address the data storage challenge of HL-LHC
  - Tape-driven workflow, where jobs get inputs from tape directly.
  - In production since 2020, for major ATLAS production campaigns like RAW data reprocessing, derivation and Monte Carlo simulation.

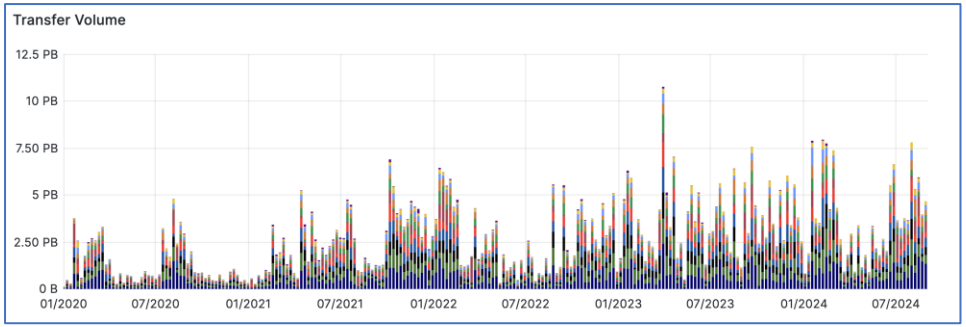

Data volume recalled from T0/T1s tape since 2020 (ATLAS DDM dashboard, bin size is weekly). Different colors represent different WLCG Clouds (~countries/regions) as the data source. From the ATLAS DDM dashboard.

DAOD on Demand ...

Overview

- Objectives
  - Large volume of DAOD datasets kept on disk but many of them lie unused
    - DAOD --- inputs to user analysis jobs
  - Can we delete them or move them to tape, then recreating them on demand ?
    - Disk space saving
    - Reduce operational load on both DDM team and physics groups
- Two scenarios to evaluate
  - Delete these DAOD datasets from disk, recreate them on demand by re-running original jobs (likely requires recalling AOD datasets from tape as inputs)
  - Move these DAOD datasets from disk to tape, recall them back using Data Carousel when needed
- Metrics
  - Time To Completion (TTC) of different scenarios - how quickly can DAOD datasets be made available to users ?
    - TTC is defined as from the time when user submits request to the time when DAOD dataset is available for user access (including steps like jobs/tasks submission, Rucio/FTS requests queuing, tape recalling, jobs running, files transferring to final destinations etc)
  - Extra load on CPU and storage resources each scenario may require.

The tests

- Data sample --- taken from DAOD exception lists of 2023 ATLAS lifetime model campaigns, excluding the following datasets:
  - Very old ones
  - Unknown to Rucio
  - Non-production related ones (e.g. validation samples)
  - Dataset size <50GB
- Recreating DAOD datasets by rerunning original jobs/tasks using PanDA
  - PanDA-- ATLAS Production and Distributed Analysis system
  - Job inputs (AOD datasets) already on disk in this test
- Tape tests
  - Tested multiple recall modes
    - single dataset recalls (DAOD and/or AOD datasets)
    - DAOD and its parent AOD dataset pair recalls
    - bulk mode recalls
  - Covered conditions with different amount of other parallel recall requests
  - Two ATLAS Tier-1 sites (FZK and RAL) participated

Results (1/2)

- Comparable TTC among different scenarios
  - <TTC of recalling DAOD from tape> = 13.1 ± 5.6 hours
  - <TTC of recalling AOD from tape> = 12.4 ± 7.2 hours
  - <TTC of recreating DAOD by jobs> = 7 ± 3 hours
- No correlation between dataset size and TTC (at individual dataset level)
- Good recall performance in bulk recall test at FZK
  - File grouping on tape (so called "smart writing") relieves concerns of suboptimal tape usage with (smaller) DAOD datasets
  - More details in the next "tape smart writing" section

No correlation between TTC and dataset size (DAOD/AOD)

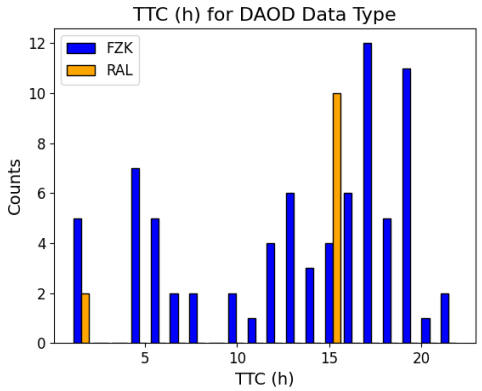

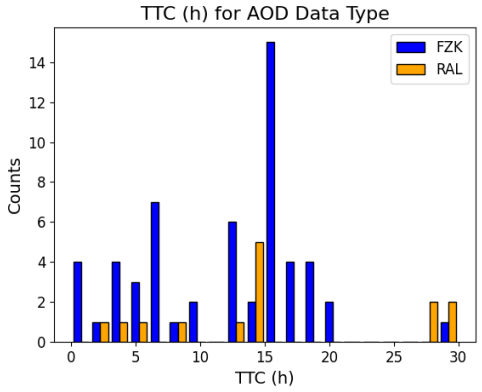

Results (2/2)

- Extra load on CPU and storage resources by different scenarios – negligible !
  - Studied DAOD access patterns in recent DAOD lifetime model exception lists
  - DAOD volume in the exception list is growing, but rarely used afterwards (plots to the right)

- Preliminary conclusion
  - All scenarios are possible and can be implemented
  - Based on the very minimal access pattern of these DAOD datasets and the comparable TTC among all scenarios, "deleting DAOD then recreating them by rerunning jobs" appears to be a feasible option.

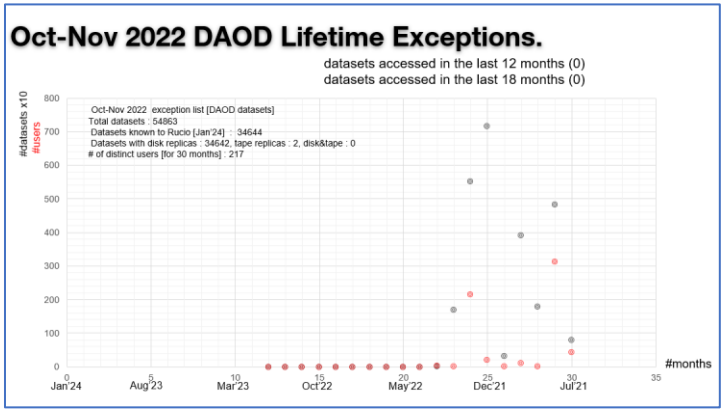

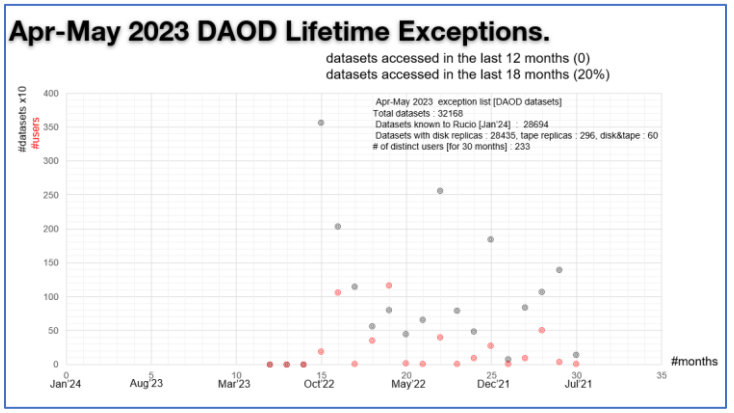

Access pattern to DAOD datasets after they are put into lifetime model exceptions list

Tape Smart Writing ...

Tape Smart Writing

- ATLAS Data Carousel represents a fundamental shift in tape usage with frequent recalls, moving away from traditional "archival-only" model. Necessitates significant changes in data writing to and reading from tape, at both ADC level and across ATLAS tape sites (mainly Tier-0/Tier-1s).

- "Smart writing" is a key strategy in Data Carousel to optimize tape bandwidth utilization
  - efficiently grouping files on tape that will be recalled together.
  - Tape recall campaigns show sites that group files on tape deliver much better performance than sites that don't, on a per tape drive basis.

Aim to put our data on tape like a well-organized warehouse

  - A "catch-all" phrase, encompassing various techniques for optimizing data layout on tape to improve read performance.
  - Reading should match how data is written on tape, the other side of the same coin, although we don't call it "smart reading".

- In this demonstrator, we run exercises with selected sites, to assess existing smart writing solutions.

Smart writing exercise with FZK Tier-1 (1/2)

- Quick overview of FZK implementation on tape smart writing
  - FZK runs dCache storage system, tape system switched from TSM to HPSS in 2023
  - In tape writing, dataset grouped together using tape file families
    - Their file families treated just like integers, not associated with specific namespaces as traditionally done
  - Files are written to tapes in file aggregates. When recalled an entire aggregate is staged (Full Aggregation Recall Mechanism)
  - 1 tape drive per file family for writing to HPSS, >40TB datasets get assigned to multiple file families/tape drives.
  - For more details, please refer to FZK presentation at CHEP-2023.
- Test procedure
  - Write data sample into FZK TSM and HPSS tape systems separately, recall back, compare performance.
  - Metric: tape bandwidth utilization on a per tape drive basis
  - Clean test environment -- minimize other parallel recall requests (e.g. those from production Data Carousel) as much as possible
  - Monitoring: both FZK internal site monitoring and ATLAS DDM dashboard.
- Data sample
  - Three data types : RAW data (RAW), reconstructed data (AOD) and derived AOD data (DAOD), ~100TB for each.
  - Size distribution of datasets in the samples : tried to mirror that in real data.
  - Specifically selected 1 or 2 >40TB datasets in the samples
- Archival metadata for grouping files on tape
  - FZK relies on archival metadata (e.g. size of dataset) to adjust the number of tape drives to use for writing a dataset to tape
  - A temporary solution, using URL parameters in transfer request, provided by Rucio to support FZK.
  - A long term solution is being tested (for details, please refer to Julien Leduc's CHEP talk)

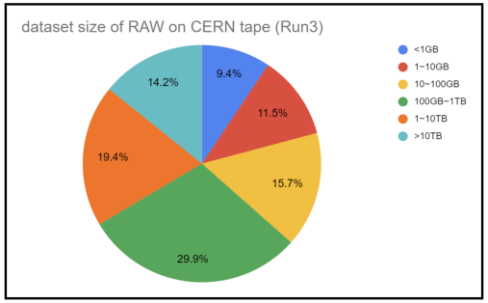

RAW dataset size distribution (based on Run3 RAW datasets on CERN tape)

Smart writing exercise with FZK Tier-1 (2/2)

- Results and observations
  - Reached 80% or higher tape bandwidth utilization for all data types (plots to the right) in HPSS, which implements smart writing mechanism, a factor of two improvements over the old TSM system.
    - Tape bandwidth utilization is defined as the ratio of delivered rate over the nominal rate per tape drive.
  - More efficient recalls at individual dataset level
    - Reduce probabilities of long tail effect (few files in a dataset take long time to recall because of being on unpopular tapes)

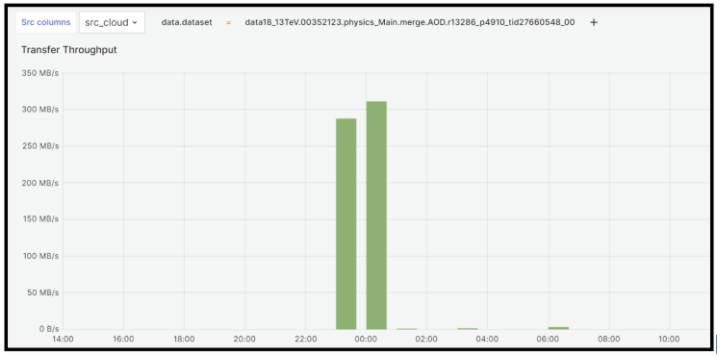

Transfer rate for a 2TB/295files AOD dataset (from ATLAS DDM dashboard)

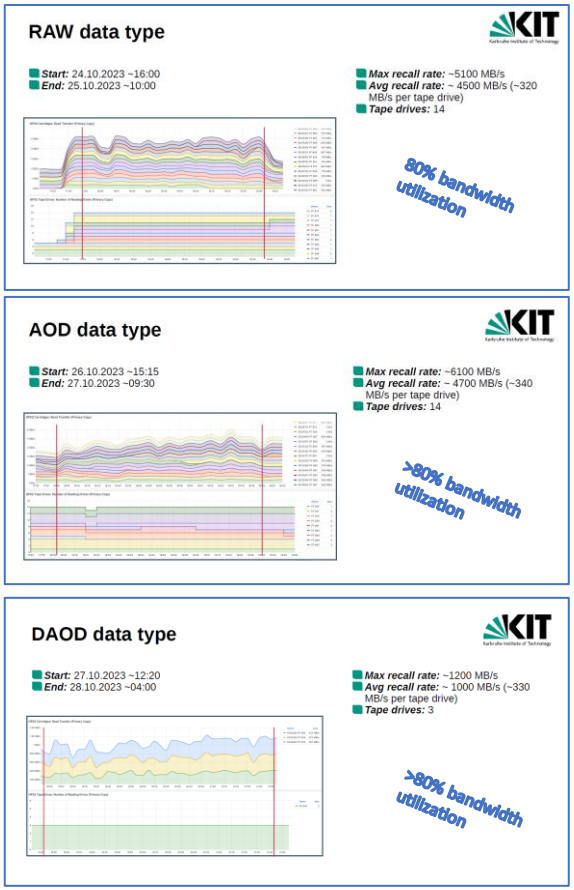

Plots from FZK internal tape monitoring, showing tape recall rate vs number of tape drives in use, for RAW, AOD and DAOD datasets respectively.

Summary and next steps

- Both demonstrator projects have achieved substantial progress, outcomes are promising.
- The DAOD-on-demand demonstrator introduces a new way of handling DAOD from lifetime mode exception list, potentially reducing disk space usage and operational overhead. Next step will be to present the preliminary results to the physics community for further feedback and evaluation.
- The tape smart writing project has demonstrated significant improvements on tape bandwidth utilization with our first selected site (FZK). For next steps:
  - Collaborate with more sites to assess their solutions
  - Continue to work with Rucio/FTS/CTA/dCache etc experts on a generic solution of providing archival metadata to sites
  - Within ADC, we will continue to evaluate and enhance our current tape workflows, in both reading and writing.